

Claims

1. A method for storing data, comprising:
transmitting a storage operation request to one of at least two controllers, the at
5 least two controllers being capable of managing communication with a plurality of
targets;
directing the storage operation request to an operational one of the at least two
controllers when the one of the at least two controllers is inoperable; and
processing the storage operation request with the operational one of the at least
10 two controllers.

2. A method for storing data as recited in claim 1, further comprising:
communicating data for storage operations using to the at least two controllers
through a network fabric interconnect.

15

3. A method for storing data as recited in claim 1, further comprising:
communicating data between at least one storage device and at least one host
using the operational one of the at least two controllers.

20 4. A method for storing data as recited in claim 1, wherein the storage
operation request is one of a request to store of data to a target and a request to read data
from a target device, the target device being a storage device.

5. A method for storing data as recited in claim 1, wherein the directing of the storage operation request includes routing the storage operation request through an L4 router.

5 6. A method for storing data as recited in claim 5, wherein the transmitting the storage request includes communicating data between a host and at least two storage devices through one of at least two L4 routers and one of the at least two controllers wherein an operable one of the at least two L4 routers is used when one of the at least two L4 routers is inoperable, the at least two L4 routers being capable of facilitating remote
10 direct memory access (RDMA) communications between the at least two target devices and the host wherein the router uses information at a transport layer to route data between transport sessions.

7. A method for storing data as recited in claim 2, wherein the network fabric
15 interconnect is an InfiniBand-type fabric, and the at least two controllers are RAID controllers.

8. A method for storing data as recited in claim 7, wherein the transmitting the storage operation request includes transmitting the request through an operational one
20 of at least two bridge chips.

9. A method for storing data as recited in claim 8, wherein the at least two bridge chips are SATA-IB chips.

10. A method for storing data as recited in claim 2, wherein the directing the storage operation request includes determining a correct path through the network fabric interconnect to a proper storage device.

5 11. A method for storing data, comprising:
providing a plurality of storage containers, each one of the plurality of storage containers having a plurality of storage devices;
generating a plurality of storage volumes where each one of the storage volumes includes at least one storage device from each one of the plurality of storage containers;
10 managing each of the plurality of volumes with a corresponding storage device controller; and
when the corresponding storage device controller is inoperable,
accessing data on the plurality of storage volumes through at least one operable storage device controller that is configured to access the volume
15 managed by the inoperable storage device controller.

12. A method for storing data as recited in claim 11, further comprising:
assigning control of each of the plurality of storage volumes to a corresponding
separate controller.

20 13. A method for storing data as recited in claim 11, further comprising:
when a controller for a particular storage container fails,
accessing data located on other storage containers through XOR
operations.

14. A method for storing data as recited in claim 11, further comprising:
managing the plurality of storage volumes by spreading the processing of
input/output requests among all of a plurality of controllers, each of the plurality of
5 controllers controlling a corresponding one of the plurality of storage volumes.

15. A method for storing data as recited in claim 11, wherein the storage
containers are RAID devices.

10 16. A method for storing data as recited in claim 11, wherein the storage
devices are disk drives.

17. A method for transmitting data in a data storage system with at least two
RAID controllers and at least two L4 routers, comprising:

15 determining functionality of the at least two L4 routers;

during a read operation,

communicating the data from a storage device to a functional L4 router;

determining at least one destination host for the data,

transferring the data to the at least one destination host using L4 routing;

20 and

during a write operation

communicating the data from a host to a functional L4 router,

determining at least one destination storage device for the data,

transferring the data to the destination storage device using L4 routing.

18. A method for transmitting data in a data storage system as recited in claim 17, wherein communicating the data from a storage device to a functional L4 router and
5 communicating the data from the host to the functional L4 router includes transmitting the data through an SATA-IB bridge and a network fabric.

19. A method for transmitting data in a data storage system as recited in claim 17, wherein the network fabric is an InfiniBand-type fabric.
10

20. A method for transmitting data in a data storage system as recited in claim 17, wherein the L4 routing includes using RDMA data transfer.

21. A storage network architecture, comprising:
15 at least two target devices;
at least two controllers for managing the at least two target devices, each of the at least two controllers configured to be capable of managing the at least two target devices when one of the at least two controllers is inoperable;

at least two switches connecting the at least two controllers and the at least two
20 target devices; and

at least two L4 routers, each one of the at least two L4 routers being capable of communicating data between a host and the at least two target devices through one of the

at least two switches and one of the at least two controllers, the L4 router being capable of facilitating remote direct memory access (RDMA) communications between the at least two target devices and the host wherein the router uses information at a transport layer to route data between transport sessions.

5

22. A network architecture as recited in claim 21, wherein each one of the at least two L4 routers utilizes at least one internal queue pair to communicate with at least one external queue pair of the host.

10

23. A network architecture as recited in claim 21, wherein the each one of the at least two L4 routers communicates with the host through reliable connection (RC) sessions.

15

24. A network architecture as recited in claim 21, wherein the each one of the at least two L4 routers communicates with the at least two controllers through reliable connection (RC) sessions.

20

25. A network architecture as recited in claim 21, wherein the router communicates with the at least two controllers through reliable datagram (RD) sessions.

26. A network architecture as recited in claim 21, wherein each one of the at least two routers communicates with the at least two target devices through reliable connection (RC) sessions.

27. A network architecture as recited in claim 21, wherein each one of the at least two L4 routers communicates with at least two target devices through reliable datagram (RD) sessions.

5

28. A network architecture as recited in claim 21, wherein the at least two L4 routers communicates with the at least two controllers through reliable datagram (RD) sessions.

10 29. A network architecture as recited in claim 21, wherein the router communicates with the at least two controllers through reliable connection (RC) sessions.

30. A network architecture as recited in claim 21, wherein the at least two controllers are RAID controllers.

15

31. A network architecture as recited in claim 21, wherein the plurality of targets is a plurality of storage devices.

32. A network architecture as recited in claim 21, wherein the plurality of
20 storage devices is a plurality of disk drives.

33. A network architecture as recited in claim 21, wherein when one of the at least two L4 routers is inoperable, the operable one of the least two L4 routers is capable of communicating data between the host and at least two target devices.

5 34. A network architecture as recited in claim 21, wherein when one of the at least two switches is inoperable, the operable one of the least two switches is capable of communicating data between the host and at least two target devices.